

Postel home

RFC 802: The ARPANET 1822L Host Access Protocol

Andrew G. Malis
Netmail: malis@bbn-unix

Bolt Beranek and Newman Inc.

November 1981

Table of Contents

1	INTRODUCTION.....	1
2	THE ARPANET 1822L HOST ACCESS PROTOCOL.....	4
2.1	Addresses and Names.....	6
2.2	Name Authorization and Effectiveness.....	8
2.3	Uncontrolled Messages.....	14
2.4	The Short-Blocking Feature.....	15
2.4.1	Host Blocking.....	16
2.4.2	Reasons for Host Blockage.....	19
2.5	Establishing Host-IMP Communications.....	22
3	1822L LEADER FORMATS.....	25
3.1	Host-to-IMP 1822L Leader Format.....	26
3.2	IMP-to-Host 1822L Leader Format.....	34
4	REFERENCES.....	42

There is one major restriction to the new capabilities being described. Both the 1822L protocol and the short-blocking feature will be implemented on C/30 IMPs only, and will therefore only be useable by hosts connected to C/30 IMPs, as the Honeywell and Pluribus IMPs do not have sufficient memory to hold the new programs and tables. This restriction also means that logical addressing cannot be used to address a host on a non-C/30 IMP. However, the ARPANET will shortly be completely converted to C/30 IMPs, and at that time this restriction will no longer be a problem.

I will try to keep my terminology consistent with that used in Report 1822, and will define new terms when they are first used. Of course, familiarity with Report 1822 (section 3 in particular) is assumed.

This document makes many references to Report 1822. As a convenient abbreviation, I will use "see 1822(x)" instead of "please refer to Report 1822, section x, for further details".

This document is a proposal, not a description of an implemented system. Thus, described features are subject to change based upon responses to this document and restrictions that become evident during implementation. However, any such changes are expected to be minor. A new RFC will be made available once the

implementation is complete containing the actual as-implemented description.

Finally, I would like to thank Dr. Eric C. Rosen, who wrote most of section 2.4, and James G. Herman, Dr. Paul J. Santos Jr., John F. Haverty, and Robert M. Hinden, all of BBN, who contributed many of the ideas found herein.

2 THE ARPANET 1822L HOST ACCESS PROTOCOL

The ARPANET 1822L Host Access Protocol, which replaces the ARPANET 1822 Host Access Protocol described in Report 1822, sections 3.3 and 3.4, allows a host to use logical addressing to communicate with other hosts on the ARPANET. Basically, logical addressing allows hosts to refer to each other using an 1822L name (see section 2.1) which is independent of a host's physical location in the network. IEN 183 (also published as BBN Report 4473) [2] gives the use of logical addressing considerable justification. Among the advantages it cites are:

- o The ability to refer to each host on the network by a name independent of its location on the network.
- o Allowing different hosts to share the same host port on a time-division basis.
- o Allowing a host to use multi-homing (where a single host uses more than one port to communicate with the network).
- o And allowing several hosts that provide the same service to share the same name.

The main differences between the 1822 and 1822L protocols are the format of the leaders that are used to introduce messages between

a host and an IMP, and the specification in those leaders of the source and/or destination host(s). Hosts have the choice of using the 1822 or the 1822L protocol. When a host comes up on an IMP, it declares itself to be an 1822 host or an 1822L host by the type of NOP message (see section 3.1) it uses. Once up, hosts can switch from one protocol to the other by issuing an appropriate NOP. Hosts that do not use the 1822L protocol will still be addressable by and can communicate with hosts that do, and vice-versa.

Another difference between the two protocols is that the 1822 leaders are symmetric, while the 1822L leaders are not. The term symmetric means that in the 1822 protocol, the exact same leader format is used for messages in both directions between the hosts and IMPs. For example, a leader sent from a host over a cable that was looped back onto itself (via a looping plug or faulty hardware) would arrive back at the host and appear to be a legal message from a real host (the destination host of the original message). In contrast, the 1822L headers are not symmetric, and a host can detect if the connection to its IMP is looped by receiving a message with the wrong leader format. This allows the host to take appropriate action upon detection of the loop.

2.1 Addresses and Names

The 1822 protocol defines one form of host specification, and the 1822L protocol defines two additional ways to identify network hosts. These three forms are 1822 addresses, 1822L names, and 1822L addresses.

1822 addresses are the 24-bit host addresses found in 1822 leaders. They have the following format:




Figure 1. 1822 Address Format

These fields are quite large, and the ARPANET will never use more than a fraction of the available address space. 1822 addresses are used in 1822 leaders only.

1822L names are 16-bit unsigned numbers that serve as a logical identifier for one or more hosts. 1822L names have a much simpler format:




Figure 2. 1822L Name Format

The 1822L names are just 16-bit unsigned numbers, except that bits 1 and 2 are not both zeros (see below). This allows over 49,000 hosts to be specified.

1822 addresses cannot be used in 1822L leaders, but there may be a requirement for an 1822L host to be able to address a specific physical host port or IMP fake host. 1822L addresses are used for this function. 1822L addresses form a subset of the 1822L name space, and have both bits 1 and 2 off.




Figure 3. 1822L Address Format

This format gives 1822L hosts the ability to directly address hosts 0-59 at IMPs 1-255 (IMP 0 does not exist). Host numbers 60-63 are reserved for addressing the four fake hosts at each IMP.

2.2 Name Authorization and Effectiveness

Every host on a C/30 IMP, regardless of whether it is using the 1822 or 1822L protocol to access the network, will be assigned at least one 1822L name (logical address). Other 1822L hosts will use this name to address the host, wherever it may be physically located. Because of the implementation constraints mentioned in the introduction, hosts on non-C/30 IMPs cannot be assigned 1822L names. To circumvent this restriction, however, 1822L hosts can use 1822L addresses to access all other hosts on the network, no matter where they reside.

At this point, several questions arise: How are these names assigned, how do they become known to the IMPs (so that translations to physical addresses can be made), and how do the IMPs know which host is currently using a shared port? To answer each question in order:

Names are assigned by a central network administrator. When each name is created, it is assigned to a host (or a group of hosts) at one or more specific host ports. The host(s) are allowed to reside at those specific host ports, and nowhere else. If a host moves, it will keep the same name, but the administrator has to update the central database to reflect the new host port. Changes to this database are distributed to the IMPs by the Network Operations Center (NOC) at BBN. For a while, the host may be allowed to reside at either of (or both) the new and old ports. Once the correspondence between a name and one or more hosts ports where it may be used has been made official by the administrator, that name is said to be authorized. 1822L addresses, which actually refer to physical host ports, are always authorized in this sense.

Once a host has been assigned one or more names, it has to let the IMPs know where it is and what name(s) it is using. There are two cases to consider, one for 1822L hosts and another for 1822 hosts. The following discussion only pertains to hosts on C/30 IMPs.

When an IMP sees an 1822L host come up on a host port, the IMP has no way of knowing which host has just come up (several hosts may share the same port, or one host may prefer to be known by

different names at different times). This requires the host to let the IMP know what is happening before it can actually send and receive messages. This function is performed by a new host-to-IMP message, the Name Declaration Message (NDM), which lists the names that the host would like to be known by. The IMP checks its tables to see if each of the names is authorized, and sends an NDM Reply to the host saying which names in the list can be used for sending and receiving messages (i.e., which names are effective). A host can also use an NDM message to change its list of effective addresses (it can add to and delete from the list) at any time. The only constraint on the host is that any names it wishes to use can become effective only if they are authorized.

In the second case, if a host comes up on a C/30 IMP using the 1822 protocol, the IMP automatically makes the first name the IMP finds in its tables for that host become effective. Thus, even though the host is using the 1822 protocol, it can still receive messages from 1822L hosts via its 1822L name. Of course, it can also receive messages from an 1822L host via its 1822L address as well. (Remember, the distinction between 1822L names and addresses is that the addresses correspond to physical locations on the network, while the names are strictly logical identifiers). The IMPs translate between the different leaders

and send the proper leader in each case (more on this below).

The third question above has by now already been answered. When an 1822L host comes up, it uses the NDM message to tell the IMP which host it is (which names it is known by). Even if this is a shared port, the IMP knows which host is currently connected.

Whenever a host goes down, its names automatically become non-effective. When it comes back up, it has to make them effective again.

Several hosts can share the same 1822L name. If more than one of these hosts is up at the same time, any messages sent to that 1822L name will be delivered to just one of the hosts sharing that name, and a RFNM will be returned as usual. However, the sending host will not receive any indication of which host received the message, and subsequent messages to that name are not guaranteed to be sent to the same host. Typically, hosts providing exactly the same service could share the same 1822L name in this manner.

Similarly, when a host is multi-homed, the same 1822L name may refer to more than one host port (all connected to the same host). If the host is up on only one of those ports, that port will be used for all messages addressed to it. However, if the

host were up on more than one port, the message would be delivered over just one of those ports, and the subnet would choose which port to use. This port selection could change from message to message. If a host wanted to insure that certain messages were delivered to it on specific ports, these messages could use either the port's 1822L address or a specific 1822L name that referred to that port alone.

Some further details are required on communications between 1822 and 1822L hosts. Obviously, when 1822 hosts converse, or when 1822L hosts converse, no conversions between leaders and address formats are required. However, this becomes more complicated when 1822 and 1822L hosts converse with each other.

The following figure illustrates how these addressing combinations are handled, showing how each type of host can access every other type of host. There are three types of hosts: "1822 on C/30" signifies an 1822 host that is on a C/30 IMP, "1822L" signifies an 1822L host (on a C/30 IMP), and "1822 on non-C/30" signifies a host on a non-C/30 IMP (which cannot support the 1822L protocol). The table entry shows the protocol and host address format(s) that the source host can use to reach the destination host.

Source Host	Destination Host		
	1822 on C/30	1822L	1822 on non-C/30
1822 on C/30	1822	1822 (note 1)	1822
1822L	1822L, using 1822L name or address (note 2)	1822L, using 1822L name or address	1822L, using 1822L address only (note 2)
1822 on non-C/30	1822	1822 (note 1)	1822

Note 1: The message is presented to the destination host with an 1822L leader containing the 1822L addresses of the source and destination hosts. If either address cannot be encoded as an 1822L address, then the message is not delivered and an error message is sent to the source host.

Note 2: The message is presented to the destination host with an 1822 leader containing the 1822 address of the source host.

Figure 4. Communications between different host types

2.3 Uncontrolled Messages

Uncontrolled messages (see 1822(3.6)) present a unique problem for the 1822L protocol. Uncontrolled messages use none of the normal ordering and error-control mechanisms in the IMP, and do not use the normal subnetwork connection facilities. As a result, uncontrolled messages need to carry all of their overhead with them, including source and destination addresses. If 1822L addresses are used when sending an uncontrolled message, additional information is now required by the subnetwork when the message is transferred to the destination IMP. This means that less host-to-host data can be contained in the message than is possible between 1822 hosts.

Uncontrolled messages that are sent between 1822 hosts may contain not more than 991 bits of data. Uncontrolled messages that are sent to and/or from 1822L hosts are limited to 32 bits less, or not more than 959 bits. Messages that exceed this length will result in an error indication to the host, and the message will not be sent. This error indication represents an enhancement to the previous level of service provided by the IMP, which would simply discard an overly long uncontrolled message without notification.

Other enhancements that are provided for uncontrolled message service are a notification to the host of any message-related errors that are detected by the host's IMP when it receives the message. A host will be notified if an uncontrolled message contains an error in the 1822L name specification, such as the name not being authorized or effective, or if the remote host is unreachable (which is indicated by none of its names being effective), or if network congestion control throttled the message before it left the source IMP. The host will not be notified if the uncontrolled message was lost for some reason once it was transmitted by the source IMP.

2.4 The Short-Blocking Feature

The short-blocking feature of the 1822 and 1822L protocols is designed to allow a host to present messages to the IMP without causing the IMP to not accept further messages from the host for long amounts of time (up to 15 seconds). It is a replacement for the non-blocking host interface described in 1822(3.7), and that description should be ignored.

2.4.1 Host Blocking

Most commonly, when a source host submits a message to an IMP, the IMP immediately processes that message and sends it on its way to its destination host. Sometimes, however, the IMP is not able to process the message immediately. Processing a message requires a significant number of resources, and when the network is heavily loaded, there can sometimes be a long delay before the necessary resources become available. In such cases, the IMP must make a decision as to what to do while it is attempting to gather the resources.

One possibility is for the IMP to stop accepting messages from the source host until it has gathered the resources needed to process the message just submitted. This strategy is known as blocking the host, and is basically the strategy that has been used in the ARPANET up to the present. When a host submits a message to an IMP, all further transmissions from that host to that IMP are blocked until the message can be processed.

It is important to note, however, that not all messages require the same set of resources in order to be processed by the IMP. The particular set of resources needed will depend on the message type, the message length, and the destination host of the message (see below). Therefore, although it might take a long time to

gather the resources needed to process some particular message, it might take only a short time to gather the resources needed to process some other message. This fact exposes a significant disadvantage in the strategy of blocking the host. A host which is blocked may have many other messages to submit which, if only they could be submitted, could be processed immediately. It is "unfair" for the IMP to refuse to accept these message until it has gathered the resources for some other, unrelated message. Why should messages for which the IMP has plenty of resources be delayed for an arbitrarily long amount of time just because the IMP lacks the resources needed for some other message?

A simple way to alleviate the problem would be to place a limit on the amount of time during which a host can be blocked. This amount of time should be long enough so that, in most circumstances, the IMP will be able to gather the resources needed to process the message within the given time period. If, however, the resources cannot be gathered in this period of time, the IMP will flush the message, sending a reply to the source host indicating that the message was not processed, and specifying the reason that it could not be processed. However, the resource gathering process would continue. The intention is that the host resubmit the message in a short time, when, hopefully, the resource gathering process has concluded

successfully. In the meantime, the host can submit other messages, which may be processed sooner. This strategy does not eliminate the phenomenon of host blocking, but only limits the time during which a host is blocked. This shorter time limit will generally fall somewhere in the range of 100 milliseconds to 2 seconds, with its value possibly depending on the reason for the blocking.

Note, however, that there is a disadvantage to having short blocking times. Let us say that the IMP accepts a message if it has all the resources needed to process it. The ARPANET provides a sequential delivery service, whereby messages with the same priority, source host, and destination host are delivered to the destination host in the same order as they are accepted from the source host. With short blocking times, however, the order in which the IMP accepts messages from the source host need not be the same as the order in which the source host originally submitted the messages. Since the two data streams (one in each direction) between the host and the IMP are not synchronized, the host may not receive the reply to a rejected message before it submits subsequent messages of the same priority for the same destination host. If a subsequent message is accepted, the order of acceptance differs from the order of original submission, and the ARPANET will not provide the same type of sequential delivery

that it has in the past.

Up to now, type 0 (regular) messages have only had sub-types available to request the standard blocking timeout. The short-blocking feature makes available new sub-types that allow the host to request messages to be short-blocking, i.e. only cause the host to be blocked for a short amount of time if the message cannot be immediately processed. See section 3.1 for a complete list of the available sub-types.

If sequential delivery by the subnet is a strict requirement, as would be the case for messages produced by NCP, the short-blocking feature cannot be used. For messages produced by TCP, however, the use of the short-blocking feature is allowed and recommended.

2.4.2 Reasons for Host Blockage

There are a number of reasons why a message could cause a long blockage in the IMP, which would result in the rejection of a short-blocking message. The IMP signals this rejection of a short-blocking message by using the Incomplete Transmission (Type 9) message, using the sub-type field to indicate which of the above reasons caused the rejection of the message. See section

3.2 for a summary of the Incomplete Transmission message and a complete list of its sub-types. The sub-types that apply to the short-blocking feature are:

6. Connection setup-delay: Although the IMP presents a simple message-at-a-time interface to the host, it provides an internal connection-oriented (virtual circuit) service, except in the case of uncontrolled messages (see section 2.3). Two messages are considered to be on the same connection if they have the same source host (i.e., they are submitted to the same IMP over the same host interface), the same priority, and the same destination host name or address. The subnet maintains internal connection set-up and tear-down procedures. Connections are set up as needed, and are torn down only after a period of inactivity. Occasionally, network congestion or resource shortage will cause a lengthy delay in connection set-up. During this period, no messages for that connection can be accepted, but other messages can be accepted.

7. End-to-end flow control: For every message that a host submits to an IMP (except uncontrolled messages) the IMP eventually returns a reply to the host indicating the disposition of the message. Between the time that the

message is submitted and the time the host receives the reply, the message is said to be outstanding. The ARPANET allows only eight outstanding messages on any given connection. If there are eight outstanding messages on a given connection, and a ninth is submitted, it cannot be accepted. If a message is refused because its connection is blocked due to flow control, messages on other connections can still be accepted.

End-to-end flow control is the most common cause of host blocking in the ARPANET at present.

8. Destination IMP buffer space shortage: If the host submits a message of more than 1008 bits (exclusive of the 96-bit leader), buffer space at the destination IMP must be reserved before the message can be accepted. Buffer space at the destination IMP is always reserved on a per-connection basis. If the destination IMP is heavily loaded, there may be a lengthy wait for the buffer space; this is another common cause of blocking in the present ARPANET. Messages are rejected for this reason based on their length and connection; messages of 1008 or fewer bits or messages for other connections may still be acceptable.

9. Congestion control: A message may be refused for reasons of congestion control if the path via the intermediate IMPs and lines to the destination IMP is too heavily loaded to handle additional traffic. Messages to other destinations may be acceptable, however.
10. Local resource shortage: Sometimes the source IMP itself is short of buffer space, table entries, or some other resource that it needs to accept a message. Unlike the other reasons for message rejection, this resource shortage will affect all messages equally, except for uncontrolled messages. The message's size or connection is not relevant.

The short-blocking feature is available to all hosts on C/30 IMPs, whether they are using the 1822 or 1822L protocol, through the use of Type 0, sub-type 1 and 2 messages. A host using these sub-types should be prepared to correctly handle Incomplete Transmission messages from the IMP.

2.5 Establishing Host-IMP Communications

When a host comes up on an IMP, or after there has been a break in the communications between the host and its IMP (see 1822(3.2)), the orderly flow of messages between the host and the

IMP needs to be properly (re)established. This allows the IMP and host to recover from most any failure in the other or in their communications path, including a break in mid-message.

The first messages that a host should send to its IMP are three NOP messages. Three messages are required to insure that at least one message will be properly read by the IMP (the first NOP could be concatenated to a previous message if communications had been broken in mid-stream, and the third provides redundancy for the second). These NOPs serve several functions: they synchronize the IMP with the host, they tell the IMP how much padding the host requires between the message leader and its body, and they also tell the IMP whether the host will be using 1822 or 1822L leaders.

Similarly, the IMP will send three NOPs to the host when it detects that the host has come up. Actually, the IMP will send six NOPs, alternating three 1822 NOPs with three 1822L NOPs. Thus, the host will see three NOPs no matter which protocol it is using. The NOPs will be followed by two Interface Reset messages, one of each style. If the IMP receives a NOP from the host while the above sequence is occurring, the IMP will only send the remainder of the NOPs and the Interface Reset in the proper style. The 1822 NOPs will contain the 1822 address of the

host interface, and the 1822L NOPs will contain the corresponding 1822L address.

Once the IMP and the host have sent each other the above messages, regular communications can commence. See 1822(3.2) for further details concerning the ready line, host tardiness, and other issues.

3 1822L LEADER FORMATS

The following sections describe the formats of the leaders that precede messages between an 1822L host and its IMP. They were designed to be as compatible with the 1822 leaders as possible. The second, fifth, and sixth words are identical in the two leaders, and all of the existing functionality of the 1822 leaders has been retained. The first difference one will note is in the first word. The 1822 New Format Flag is now also used to identify the two types of 1822L leaders, and the Handling Type has been moved to the second byte. The third and fourth words contain the Source and Destination 1822L Name, respectively.

3.1 Host-to-IMP 1822L Leader Format




Figure 5. Host-to-IMP 1822L Leader Format

Bits 1-4: Unused, must be set to zero.

Bits 5-8: 1822L Host-to-IMP Flag:

This field is set to decimal 13 (1101 in binary).

Bits 9-16: Handling Type:

This field is bit-coded to indicate the transmission characteristics of the connection desired by the host. See 1822(3.3).

Bit 9: Priority Bit:

Messages with this bit on will be treated as priority messages.

Bits 10-16: Unused, must be zero.

Bits 17-20: Unused, must be zero.

Bit 21: Trace Bit:

If equal to one, this message is designated for tracing as it proceeds through the network. See 1822(5.5).

Bits 22-24: Leader Flags:

Bit 22: A flag available for use by the destination host.

See 1822(3.3) for a description of its use by the IMP's TTY fake host.

Bits 23-24: Reserved for future use, must be zero.

Bits 25-32: Message Type:

Type 0: Regular Message - All host-to-host communication occurs via regular messages, which have several sub-types, found in bits 77-80. These sub-types are:

0: Standard - The IMP uses its full message and error control facilities, and host blocking (see section 2.4) may occur.

1: Standard, short-blocking - See section 2.4.

2: Uncontrolled, short-blocking - See section 2.4.

3: Uncontrolled - The IMP will perform no message-control functions for this type of message, and network flow and congestion control (see section 2.4) may cause loss of the message. Also see 1822(3.6) and section 2.3.

4-15: Unassigned.

Type 1: Error Without Message ID - See 1822(3.3).

Type 2: Host Going Down - see 1822(3.3).

Type 3: Name Declaration Message (NDM) - This message is used by the host to declare which of its 1822L names is or is not effective (see section 2.2), or to make all of its names non-effective. The first 16 bits of the data portion of the NDM message, following the leader and any padding, contains the number of 1822L name

entries contained in the message. This is followed by the 1822L name entries, each 32 bits long, of which the first 16 bits is a 1822L name and the second 16 bits contains either of the integers zero or one. Zero indicates that the name should not be effective, and one indicates that the name should be effective. The IMP will reply with a NDM Reply message (see section 3.2) indicating which of the names are now effective and which are not. Pictorially, a NDM message has the following format (including the leader, which is printed in hexadecimal):




Figure 6. NDM Message Format

An NDM with zero entries will cause all current effective names for the host to become non-effective.

Type 4: NOP - This allows the IMP to know which style of leader the host wishes to use. A 1822L NOP signifies that the host wishes to use 1822L leaders, and an 1822 NOP signifies that the host wishes to use 1822 leaders. All of the other remarks concerning the NOP message in

1822(3.3) still hold. The host should always issue NOPs in groups of three to insure proper reception by the IMP. Also see section 2.5 for a further discussion on the use of the NOP message.

Type 8: Error with Message ID - see 1822(3.3).

Types 5-7,9-255: Unassigned.

Bits 33-48: Source Host:

This field contains one of the source host's 1822L names (or, alternatively, the 1822L address of the host port the message is being sent over). This field is not automatically filled in by the IMP, as in the 1822 protocol, because the host may be known by several names and may wish to use a particular name as the source of this message. All messages from the same host need not use the same name in this field. Each source name, when used, is checked for authorization, effectiveness, and actually belonging to this host. Messages using names that do not satisfy all of these requirements will not be delivered, and will instead result in an error message being sent back into the source host. If the host places its 1822L Address in this field, the address is checked to insure that it actually represents the host port where the message originated. If the message is destined for an 1822 host on a non-C/30 IMP, this field MUST

contain the source host's 1822L address (see Figure 4 in section 2.2).

Bits 49-64: Destination Host:

This field contains the 1822L name or address of the destination host. If it contains a name, the name will be checked for effectiveness, with an error message returned to the source host if the name is not effective. If the message is destined for an 1822 host on a non-C/30 IMP, this field MUST contain the destination host's 1822L address (see Figure 4 in section 2.2).

Bits 65-76: Message ID:

This is a host-specified identification used in all type 0 and type 8 messages, and is also used in type 2 messages. When used in type 0 messages, bits 65-72 are also known as the Link Field, and should contain values specified in Assigned Numbers [3] appropriate for the host-to-host protocol being used.

Bits 77-80: Sub-type:

This field is used as a modifier by message types 0, 2, 4, and 8.

Bits 81-96: Unused, must be zero.

3.2 IMP-to-Host 1822L Leader Format




Figure 7. IMP-to-Host 1822L Leader Format

Bits 1-4: Unused and set to zero.

Bits 5-8: 1822L IMP-to-Host Flag:

This field is set to decimal 14 (1110 in binary).

Bits 9-16: Handling Type:

This has the value assigned by the source host (see section 3.1). This field is only used in message types 0, 5-9, 11 and 15.

Bits 17-20: Unused and set to zero.

Bit 21: Trace Bit:

If equal to one, the source host designated this message for tracing as it proceeds through the network. See 1822(5.5).

Bits 22-24: Leader Flags:

Bit 22: Available as a destination host flag.

Bits 23-24: Reserved for future use, set to zero.

Bits 25-32: Message Type:

Type 0: Regular Message - All host-to-host communication occurs via regular messages, which have several sub-types. The sub-type field (bits 77-80) is the same as sent in the host-to-IMP leader (see section 3.1).

Type 1: Error in Leader - See 1822(3.4).

Type 2: IMP Going Down - See 1822(3.4).

Type 3: NDM Reply - This is a reply to the NDM host-to-IMP message (see section 3.1). It will have the same number of entries as the NDM message that is being replying to, and each listed 1822L name will be accompanied by a zero or a one. A zero signifies that the name is not effective, and a one means that the name is now effective.

Type 4: NOP - The host should discard this message. It is used during initialization of the IMP/host communication. The Destination Host field will contain the 1822L Address of the host port over which the NOP is being sent. All other fields are unused.

Type 5: Ready for Next Message (RFNM) - See 1822(3.4).

Type 6: Dead Host Status - See 1822(3.4).

Type 7: Destination Host or IMP Dead (or unknown) - This message is sent in response to a message for a destination which the IMP cannot reach. The message to the "dead" destination is discarded. See 1822(3.4) for a complete list of the applicable sub-types. If this message is in response to a standard (type 0, sub-type 0 or 1) message, it will be followed by a Dead Host Status message, which gives further information about

the status of the dead host. If this message is in response to an uncontrolled (type 0, sub-type 2 or 3) message, only sub-type 1 (The destination host is not up) will be used, and it will not be followed by a Dead Host Status message.

Type 8: Error in Data - See 1822(3.4).

Type 9: Incomplete Transmission - The transmission of the named message was incomplete for some reason. An incomplete transmission message is similar to a RFNM, but is a failure indication rather than a success indication. This message is also used by the short-blocking feature to indicate that the named message was rejected because it would have caused to IMP to block the host for a long amount of time. See section 2.4 for more details concerning the short-blocking feature.

The message's sub-types are:

0: The destination host did not accept the message quickly enough.

1: The message was too long.

2: The host took more than 15 seconds to transmit the message to the IMP. This time is measured from the last bit of the leader through the last bit of the message.

3: The message was lost in the network due to IMP or circuit failures.

4: The IMP could not accept the entire message within 15 seconds because of unavailable resources. This sub-type is only used in response to non-short-blocking messages. If a short-blocking message timed out, it will be responded to with one of the sub-types 6-10.

5: Source IMP I/O failure occurred during receipt of this message.

Sub-types 6-10 are all issued in response to a short-blocking message that timed out (would have caused the host to become blocked for a long amount of time). The sub-types are designed to give the host some indication of why it timed out and what other messages would also time out. See section 2.4.2 for further details concerning each of these sub-types.

6: The message timed out because of connection set-up delay. Further messages to the same host (if on the same connection) may also be affected.

7: The message timed out because of end-to-end flow control. Further messages to the same host on the same connection will also be affected.

8: Destination IMP buffer shortage caused the message to time out. This affects multi-packet standard messages to the specified host, but shorter messages or messages to hosts on other IMPs may not be affected.

9: Network congestion control caused the message to be rejected. Messages to hosts on other IMPs may not be affected, however.

10: Local resource shortage kept the IMP from being able to accept the message within the short-blocking timeout period.

11-15: Unassigned.

Type 10: Interface Reset - See 1822(3.4).

Type 15: 1822L Name or Address Error - This message is sent in response to a type 0 message from a host that contained an erroneous Source Host or Destination Host field. Its sub-types are:

0: The Source Host 1822L name is not authorized or not effective.

1: The Source Host 1822L address does not match the host port used to send the message.

2: The Destination Host 1822L name is not authorized.

3: The Destination Host 1822L name is authorized but

not effective, even though the named host is up.

If the host were actually down, a type 7 message would be returned, not a type 15.

4: The Source or Destination Host field contains a 1822L name, but the host being addressed is on a non-C/30 IMP (see Figure 4 in section 2.2).

5-15: Unassigned.

Types 11-14,16-255: Unassigned.

Bits 33-48: Source Host:

For type 0 messages, this field contains the 1822L name or address of the host that originated the message. All replies to the message should be sent to the host specified herein. For message types 5-9, 11 and 15, this field contains the source host field used in a previous type 0 message sent by this host.

Bits 49-64: Destination Host:

For type 0 messages, this field contains the 1822L name or address that the message was sent to. This allows the destination host to detect how it was specified by the source host. For message types 5-9, 11 and 15, this field contains the destination host field used in a previous type 0 message sent by this host.

Bits 65-76: Message ID:

For message types 0, 5, 7-9, 11 and 15, this is the value assigned by the source host to identify the message (see section 3.1). This field is also used by message types 2 and 6.

Bits 77-80: Sub-type:

This field is used as a modifier by message types 0-2, 4-7, 9, 11 and 15.

Bits 81-96: Message Length:

This field is contained in type 0 and type 3 messages only, and is the actual length in bits of the message (exclusive of leader, leader padding, and hardware padding) as computed by the IMP.

4 REFERENCES

- [1] Specifications for the Interconnection of a Host and an IMP,
BBN Report 1822, May 1978 Revision.
- [2] E. C. Rosen et. al., ARPANET Routing Algorithm Improvements,
IEN 183 (also published as BBN Report 4473, Vol. 1), August
1980, pp. 55-107.
- [3] J. Postel, Assigned Numbers, RFC 790, September 1981, p. 10.

INDEX

1822.....	4
1822 address.....	6
1822 host.....	5
1822L.....	4
1822L address.....	7
1822L host.....	5
1822L name.....	6
authorized.....	9
blocking.....	16
congestion control.....	22, 39
connection.....	20, 38
destination host.....	32, 40
effective.....	10
flow control.....	20, 38
handing type.....	27, 35
incomplete transmission message.....	19, 37
leader flags.....	27, 35
link field.....	32
logical addressing.....	4
message ID.....	32, 41
message length.....	41
message type.....	28, 35
multi-homing.....	4
NDM.....	10, 28
NDM reply.....	10, 36
NOC.....	9
NOP.....	5, 22, 30, 36
outstanding.....	21
priority bit.....	27
regular message.....	28, 35
RFNM.....	36
short-blocking feature.....	15
short-blocking message.....	19, 28
source host.....	31, 40
standard message.....	28
sub-type.....	32, 41
symmetric.....	5
trace bit.....	27, 35
uncontrolled message.....	14, 28

